

SVD analysis for radiographic object reconstruction II: Null space enhancements

by

Tom Asaki and Kevin R. Vixie

Abstract

This report presents 2D radiographic reconstructions using the method of Prior Singular Value Decomposition (pSVD). The technique utilizes (typically minimal) prior knowledge of the object to enhance reconstructions by the inclusion of projection operator null space vectors. Examples of prior knowledge illustrated in this report are that the object density is: (1) non-negative; (2) bounded above and below; (3) a known constant; and (4) of a known set of discrete values. Several test objects and simulated *noisy* data are used to illustrate the method.

One set of actual proton radiographs taken at LANSCE is used to verify the method. The radiographed object (BCO4) was viewed from 31 equally spaced angles and remained static. The pSVD reconstructions are shown to be superior to the standard filtered back projection reconstructions.

It is found that known constraints improve reconstructions by any reasonable metric. Even for the simple constraints discussed here, the null space enhancements can be significant. In some cases sparse noisy data is sufficient for essentially exact reconstructions. The implications for experiments with limited data is obvious.

Natural extensions to 3D reconstructions and dynamic constraints are discussed. Appendices catalogue many reconstruction examples of both standard test objects and the BCO4.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 3 |
| 1.1 | Quick Review of the Problem | 3 |
| 1.2 | Null Space Enhancements and the pSVD Method | 4 |
| 2 | Examples Using Test Objects | 6 |
| 2.1 | Object Descriptions | 6 |
| 2.2 | Data Generation | 8 |
| 2.3 | Examples of Prior Knowledge | 8 |
| 2.4 | Results and Discussion | 9 |
| 2.4.1 | TO1 | 9 |
| 2.4.2 | TO2 | 10 |
| 2.4.3 | TO3 | 10 |
| 2.4.4 | TO4 | 10 |
| 2.4.5 | TO5 | 10 |
| 2.4.6 | Reconstruction Norms | 16 |
| 3 | LANSCE BCO4 – An Example With Real Data | 18 |
| 3.1 | Experiment and Data Description | 18 |
| 3.2 | Reconstructions and General Discussion | 21 |
| 3.3 | Comparative Analysis | 29 |
| 4 | AHF and the Number-of-Views Question Revisited | 34 |
| 4.1 | Fixed Resolution Study | 34 |
| 4.2 | Fixed Number-of-Views Study | 40 |
| 5 | Conclusion | 45 |
| 6 | Appendix: pSVD Reconstruction Examples | 45 |
| 7 | References | 62 |

can be linearly projected onto detectors with M total pixels. For the underdetermined problem $M < N$. The SVD of Π is illustrated in figure 1. The decomposition can be written $\Pi = Q * \Sigma * P^T$ where Q and P are orthogonal and Σ is diagonal. The columns of Q span the range of Π and thus provide an orthogonal basis for the data space. The rows of P^T span the domain of Π providing a basis for the object representation space. The singular values σ_i are non-negative and arranged in decreasing order. Let σ_k be the smallest singular value such that $\sigma_k \geq \sigma_{cut}$ for some cutoff value σ_{cut} (chosen for controlling noise amplification in inversions from the data space into the object space). Then the first k rows of P^T span the space of reconstructible objects, and the last $N - k$ rows of P^T span the space of objects invisible to Π (the null space of Π).

1.2 Null Space Enhancements and the pSVD Method

The N -dimensional object space is illustrated in two dimensions in Fig. 2. The space is divided into the null space, $\mathbf{N} = span\{p_{k+1}, p_{k+2}, \dots, p_N\}$ and the reconstruction space, $\mathbf{V} = span\{p_1, p_2, \dots, p_k\}$. The object is assumed to be sufficiently well represented by some object X within this space. (If the object is exactly describable within $span\{p_1, \dots, p_N\}$ then X is the object.) The initial SVD reconstruction is given by X_0 (the projection of X onto \mathbf{V}). The only difference between X and X_0 is some unknown null space vector, $\Pi_{\mathbf{N}}X$.

Now suppose some object Y_1 can be chosen within our space whose null space projection is directed toward X . In other words, suppose a Y_1 can be chosen based on our knowledge of X_0 and possibly some knowledge of X such that $X_1 \equiv X_0 + \Pi_{\mathbf{N}}Y_1$ is a better approximation of X than X_0 . Next choose a new object Y_2 based upon knowledge of any or all previous reconstructions. Then the new null space correction leads to X_2 . This process can be repeated indefinitely, and if the choices of Y_i are careful this leads to a much improved reconstruction.

Clearly, it is possible to make a poor choice for any Y_i , for example Y_{bad} , so that the correction null vector leads away from X . So it might seem that within the $N - k$ dimension null space a poor choice is likely. Fortunately, this turns out to not be the case for some very important and useful experiments.

The effects of noise in the data can pose some problems for this type of enhancement. The result is that the X_i will not always converge toward X but to some object \bar{X} within some neighborhood of X . The neighborhood is determined by prior knowledge of the object used in selecting the Y_i . Fortunately, again, even minimal knowledge about the object can tightly constrain the size of the neighborhood for many cases of interest.

The central problem of the Prior Singular Value Decomposition (pSVD) method is a systematic and useful selection of the Y_i . How does one choose the prior objects which give helpful null space corrections? How much knowledge about the actual object is necessary? When is it possible to obtain an essentially exact reconstruction ($X_i \rightarrow X$)? These questions are addressed by example in the following section.

It is important to note how the pSVD method treats noise in the data. The SVD of the projection matrix and careful inversion ensures that data noise is not magnified onto the reconstruction

space. That's all! Knowledge of the nature of the data is not required. Noisy data is not smoothed or otherwise altered. The success of the pSVD method, however, does depend upon the statistical nature of the noise. If the noise is Poisson and uncorrelated among detector pixels then minimal prior knowledge allows pSVD to obtain excellent reconstructions from within initially noisy reconstructions.




Figure 2: A 2D representation of the N -dimensional object space for a typical underdetermined inversion. The reconstruction space \mathbf{V} is spanned by k vectors. The null space \mathbf{N} is spanned by the remaining $N - k$ vectors. The object is given by X and the SVD reconstruction is $\Pi_{\mathbf{V}}X = X_0$. Objects used as priors for the pSVD method are shown as Y_i and the corresponding pSVD reconstructions as X_i .

2 Examples Using Test Objects

Several test experiments were conducted for testing the usefulness of the pSVD method. This section begins by describing the test objects and the data generation process for these examples. Next, three prior-knowledge scenarios are described which are commonly found in radiographic experiments. Finally, many examples of reconstructed objects are shown and discussed.

2.1 Object Descriptions

Five 2D standard test objects were used in this study of the pSVD methods. They have been previously described, but we provide a brief description here. The objects are shown in figure 3. Each object is exactly defined on analytically defined on a 2D region independent of any finite sampling. For the purpose of displaying the figure, however, each object is sampled on a 512×512 square voxel grid. The objects are known by the names TOn (test object number n). All densities lie in the range $[0,1]$.

The first object (TO1) has unit density in a region between two non-concentric circles and zero density otherwise. This object might represent, for example, a 2D transverse section of a cylindrical rod with a removed axially-parallel cylindrical section. It is exactly binary in density.

The second object (TO2) has unit density inside a circle (the same size as the outer bounding edge of TO1) except for six overlapping zero-density triangular regions. This object is an idealization of high-stress damage which might occur to a cylindrical rod. The void fraction is 10.5%.

The third object (TO3) is a sum of constant density circular patches with a circular zero density patch stamped out of the middle. Thus, the density distribution is strictly discrete with values from the set $\{0, 0.25, 0.5, 0.75, 1\}$. This object can be thought of as an idealization of a collection of uniform density objects.

The fourth object (TO4) has a smooth spatial density-distribution. The densities vary nearly continuously from zero to unity over the object. This object can be thought of as a 2D slice of a compressible material in a dynamic experiment.

The fifth object (TO5) is the average of the previous two objects. It is the most complex of the test objects. It displays a density profile that is smooth except for the discrete jumps corresponding to the circular outlines of TO3. TO5 might represent a compressible material with the presence of shocks.

Each of these objects is defined by densities independent of any finite sampling. Reconstructions, however, are given by a collection of density values on a voxelation of finite precision(?wrong word?). Thus, exact reconstructions are not possible without knowledge of the actual object-parameterizations. For the purpose of characterizing reconstruction merit, we have chosen to compare an N -voxel reconstruction with a $100N$ -voxel approximate object. Such a finely (but finitely) sampled object representation provides a balance between true object characterization and computational efficiency.




Figure 3: Finitely-sampled representations of the five test objects: (TO1) binary density offset ring; (TO2) binary density internal damage simulation; (TO3) discrete-density stacked circles; (TO4) smoothly varying density dynamic snapshot simulation; (TO5) average of TO3 and TO4.

2.2 Data Generation

The data used for object reconstructions is generated based on the following scenario. Suppose a collimated beam of particles passes through the object (the beam being in the plane of the 2D object). The particle stream along any line is attenuated exponentially according to the total mass through which it passes. In this way ‘density’ can be interpreted as an attenuation density rather than a mass density. A detector counts the particles which pass through or past the object. The detector is discretized into J pixels. Several of these 1D views of the object are obtained by considering different beam angles through the object. Together, these L views of J pixels each compose the noiseless data set of $M = JL$ numbers. For the cases considered in this section $J = 200$, $L = 12$, and the views are equally spaced at angles within 180 degrees. The total detector width is the same width as the square voxel collection defining the reconstruction (the length of J pixels is equal to the length of K voxels).

Experimental random noise is added into the data on the assumption that the particle counts should exhibit a Poisson distribution without correlation between detector pixel counts. The noise level will thus be larger for pixels corresponding to beam paths through thick object regions and smallest for pixels corresponding to beam paths passing outside of the object. The noise level for these studies (unless otherwise noted) is 5% for the lowest particle counts and 1% for the highest particle counts. Nonlinear mechanisms, such as multiple scattering and effects of beam strength, have been ignored.

The noisy particle count data is then inverted to represent a projection of the object mass onto the detector planes. This noisy data is the input by which the pSVD method attempts to reconstruct the object.

2.3 Examples of Prior Knowledge

Four examples of prior knowledge are considered along with corresponding prior-object definitions. Each example is physically reasonable and together they encompass a large body of real applications.

The first example is that objects do not have negative density. In the attenuation density sense this is equivalent to saying that the object has no particle sources. The prior object Y obtained from a reconstruction X is defined by the elements $y_n = \max(0, x_n)$ where n indexes the N reconstruction voxels.

The second example is a generalization of the first, namely that the density profile is bounded above and below. Suppose that some upper bound on the density ρ_{max} is known. Then the prior object Y obtained from a reconstruction X is defined by the elements $y_n = \max(\min(x_n, \rho_{max}), 0)$. Typically the lower bound density is zero since the experiment is best designed to encompass objects of interest. But further generalization is obvious.

The third example is that the object is composed of a material of only one nonzero density. While the second example allows for continuous densities between 0 and ρ_{max} , this example confines the densities to either the value 0 or the value ρ_{max} . In this case it is advantageous to use the prior of the second example in the pSVD method. Then this prefinal reconstruction is modified by

imposing a density cutoff value ρ_{cut} so that densities above this value are set to ρ_{max} and densities below this value are set to zero. The cutoff value is chosen so that the object retains the correct total mass (determined from the data).

The fourth example is that the object is composed of a discrete set of known densities. This is the natural extension of the third example and is treated analogously.

2.4 Results and Discussion

Each of the six test objects are compared with five reconstructions from the same noisy data set. The reconstruction methods are the standard SVD with attention to noise ($\sigma_k \equiv 1$), and enhancements using the four prior knowledge methods described earlier. Grayscale images are presented in figures 4 through 8 and various norms on the reconstructions are given together in table [ref].

For each figure, subfigure (a) is the object and the remaining subfigures are the pSVD reconstructions using the following prior information: (b) no prior information (SVD); (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. Each reconstruction includes only the voxels within a circular area enclosed by the voxel grid. This is most apparent when viewing the SVD reconstructions of subfigures (b). The shade of gray outside this circle but inside the square represents the density value zero. Darker shades indicate negative values and lighter shades indicate positive values. Each object has a maximal density of one.

It is worth repeating that all of the reconstructions shown in this section are performed on noisy data. Reconstructions from noiseless data can only be better, and in any case are of academic value only. Some examples of the use of noiseless data are examined in the appendix.

2.4.1 TO1

The first and simplest object (figure 4) is reconstructed well using a variety of priors. The initial SVD solution captures the low spatial frequency information but fails to recover the sharp boundaries of the two defining circles. The characteristic is common to reconstructions from sparse data. It is analogous to the imperfect fourier reconstruction of a step function from a finite number of basis harmonic functions. The result is a smoothed boundary and oscillations away from the edge. These features are also apparent in the SVD reconstruction.

Using the non-negativity prior (figure 4c) there is significant improvement. The boundaries are sharpened, much of the zero density region is correctly recovered, and the density oscillations within the object are much reduced. The further assumption bounding the density from above to a value of 1 yields (figure 4d) an essentially exact solution insofar as the voxelations allows. The other priors of this study do not provide additional improvement since they are based upon this bounded density solution.

2.4.2 TO2

The second test object (figure 5) is also reconstructed well using a variety of priors. The initial SVD solution once again captures the low spatial frequency information. However, it fails to accurately define the interior structure which is largely of high spatial frequency. Even the non-negative density prior (figure 5c) does not provide significant improvement, although the outer boundary is accurately described. But the bounded density prior (figure 5d) does manage to reconstruct the interior features with remarkable accuracy. With this solution as a base, the binary density prior (figure 5e) yields a solution that is nearly (but not quite!) exact. The discrete density prior (figure 5f) does not perform as well because the assumption is incorrect. Nevertheless, even this ‘poor’ assumption far outperforms the standard SVD reconstruction.

2.4.3 TO3

The third test object (figure 6) illustrates a case where a discretized density is known. The SVD solution is typical with smoothed boundaries and oscillations with spatial frequency characteristic of the number of views. Imposing non-negativity in the density (figure 6c) provides remarkable enhancement with boundary sharpening and oscillation suppression. Bounding the density (figure 6d) does not make significant additional changes since much of this object’s density lies well below the maximum density. Assuming that the density is binary (figure 6e) is a very bad prior indeed. The reconstruction retains only the most rudimentary resemblance to the actual object. This is simply a consequence of the object being not binary in density. And finally, if the discrete density set is known and imposed as a prior (figure 6f) then the object is nearly exactly recovered.

2.4.4 TO4

The fourth test object (figure 7) is smooth. This object is the most complex examined thus far. The now familiar SVD solution is not surprising as a blurred and wrinkled version of the object. The non-negative density and bounded density solutions are much improved reconstructions showing improved zero density detection and significant high-spatial frequency content. However they retain many spurious features. The final two reconstructions are based upon discrete density assumptions and perform quite poorly as expected.

2.4.5 TO5

The fifth test object (figure 7) combines the tasks of reproducing both smooth and sharp features. The SVD reconstruction does a fair job at capturing the low spatial frequency content of both. The bounded density reconstructions are much better, providing the characteristic boundary sharpening, spurious oscillation reduction, and zero density detection. Again, the final two reconstructions are based upon discrete density assumptions and perform quite poorly as expected.




Figure 4: Reconstruction examples of object TO1. This object (a) has binary density. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$




Figure 5: Reconstruction examples of object TO2. This object (a) has binary density. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$




Figure 6: Reconstruction examples of object TO3. This object (a) has five evenly-spaced discrete densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$




Figure 7: Reconstruction examples of object TO4. This object (a) has a (voxelized) smooth density distribution from zero to one. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$




Figure 8: Reconstruction examples of object TO5. This object (a) has a (voxelized) smooth density distribution from zero to one interrupted by circularly bounded discrete density jumps. This object is the average of TO3 and TO4. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$

2.4.6 Reconstruction Norms

The choice of norm (metric, function of merit, etc.) for some real experiment depends upon the features of interest. Some possible examples might include: the total absolute density deviation over the object; the maximal density deviation over the object; the length of some critical boundary; and the object volume containing a certain density. Clearly all possible norms cannot be examined for all possible cases. Only a detailed case-by-case experiment evaluation can provide the proper norm.

For this study the reconstruction merit is examined in terms of three simple norms that capture a large portion of interesting cases. These norms provide a guideline for considering the merit of various prior methods. For an object X and a reconstruction R with voxel density values X_n and R_n , the chosen norms are the normalized L_1 norm

$$L_1(R) \equiv \frac{\sum_n \text{abs}(X_n - R_n)}{\sum_n \text{abs}(X_n)}, \quad (1)$$

the normalized L_2 norm

$$L_2(R) \equiv \frac{(\sum_n (X_n - R_n)^2)^{1/2}}{(\sum_n (X_n)^2)^{1/2}}, \quad (2)$$

and the normalized L_∞ norm

$$L_\infty(R) \equiv \frac{\max_n (X_n - R_n)}{\max_n (X_n)}. \quad (3)$$

For example, the L_∞ norm is the maximal density deviation over the reconstruction normalized to the maximal density in the object.

The results for the reconstructions of the test objects are given in table 1. The left column lists the objects and the other columns indicate the method of reconstruction, SVD or type of pSVD. For each object-reconstruction pair the three numbers are the three norms listed in order from top to bottom. This table quantifies the discussion of the individual reconstructions.

The first three objects (those of discrete density profile) can be reconstructed very accurately. This is a consequence of the favorable match of object definition and choice of prior. Large values of the L_∞ norm reflect the fact that sharp density boundaries are not precisely reproduced – this norm finds the worst case. Bear in mind, however, that the reconstructions are from *noisy* data.

Smooth objects are more difficult to reconstruct and are best handled through bounded-density priors (although other better priors are possible). Objects TO4 and TO5 are nevertheless reproduced with an L_2 norm within 15%. Contrast this with the original SVD L_2 norms near 30%.

| | SVD | $\rho \geq 0$ | $0 \leq \rho \leq 1$ | $\rho \in \{0, 1\}$ | $\rho \in \{\rho_i\}$ |
|-----|-------|---------------|----------------------|---------------------|-----------------------|
| TO1 | 0.180 | 0.055 | 0.034 | 0.023 | 0.029 |
| | 0.190 | 0.128 | 0.122 | 0.151 | 0.124 |
| | 0.784 | 1.000 | 1.000 | 1.000 | 1.000 |
| TO2 | 0.236 | 0.140 | 0.052 | 0.035 | 0.045 |
| | 0.250 | 0.208 | 0.154 | 0.188 | 0.156 |
| | 1.207 | 1.153 | 1.000 | 1.000 | 1.000 |
| TO3 | 0.292 | 0.110 | 0.107 | 0.823 | 0.043 |
| | 0.216 | 0.147 | 0.146 | 0.813 | 0.150 |
| | 0.659 | 0.757 | 0.799 | 1.000 | 0.750 |
| TO4 | 0.377 | 0.195 | 0.166 | 0.390 | 0.205 |
| | 0.291 | 0.200 | 0.170 | 0.416 | 0.206 |
| | 0.543 | 0.436 | 0.434 | 0.896 | 0.507 |
| TO5 | 0.269 | 0.151 | 0.151 | 0.749 | 0.216 |
| | 0.214 | 0.158 | 0.158 | 0.722 | 0.228 |
| | 0.414 | 0.390 | 0.390 | 0.840 | 0.408 |

Table 1: Reconstruction norms for the various object-prior pairs. The three values represent the L_1 , L_2 , and L_∞ norms, respectively.

3 LANSCE BCO4 – An Example With Real Data

We now present concrete examples of SVD and pSVD reconstruction techniques using actual proton radiography data. The experiment was performed at the Los Alamos Neutron Science Center (LANSCE) facility by a team of researchers in LANSCE-1 during (month) of (year).

3.1 Experiment and Data Description

The radiographed object is known as BCO4. It is actually two distinct solid 3D objects that are statically mounted using styrofoam spacers. Multiple-view 2D proton-count data were obtained by mounting the BCO4 on a horizontal rotation table and performing 31 single-view shots separated by 6 degrees. The data were then preprocessed to eliminate spurious values and detector-tiling edge effects. The preprocessed radiograph of the BCO4 at 90 degrees is shown in Fig. 9. This radiograph, and all others taken at the various rotation angles, forms a 1169 (vertical) by 1220 (horizontal) pixelated 2D data set. The gray scale is indicative of the proton particle counts recieved at each detector pixel. Darker regions indicate lower particle counts and the lightest regions indicate the greatest particle counts. Thus the large light gray area surrounding the objects represents the background associated with a nonattenuated beam.

A second view of the BCO4 is shown in Fig. 10. This view is the zero degree view which is perpendicular to that in Fig. 9. A quick visual comparison of these two views reveals much about the objects already. For example, the central object appears to be a short cylinder with an attached wedge-shaped protrusion. Opposite this feature there appears to be a hemispherical(?) dimple removed from the cylinder. Within this structure is a dense cone-shaped object within which are two cone-shaped regions of low (zero?) density. The smaller object can be described similarly.




Figure 9: The preprocessed radiograph of the BCO4 at the view angle of 90 degrees. Shades of gray indicate particle count values with the darkest regions representing the fewest counts.




Figure 10: The preprocessed radiograph of the BCO4 at the view angle of zero degrees. Shades of gray indicate particle count values with the darkest regions representing the fewest counts.

3.2 Reconstructions and General Discussion

For testing the pSVD method we considered the 1D data set constructed by selecting a single row of data values from each of the 30 views. Our choice was to use pixel row 610 (numbered from the top) from each view. These data are the particle counts for that part of the proton beam passing through the approximate center of the large object. The 2D reconstruction from this data will be a 2D density representation of this object in a plane perpendicular to the detector planes.

The full inversion of this data set into a high resolution reconstruction requires much more RAM than is available on our present workstation (4GB). For this problem we would like to compute the SVD of a dense matrix of size mv by mv (36,600 square). Our current matlab code and available RAM limits us to $mv < 7000$, assuming $n = m$. Thus we must reduce the number of views used for the reconstruction, reduce the data set resolution, or some combination of the two. In order to gain the most from the pSVD null space correctors we consider a 30 view reconstruction at a resolution of $n = m = 220$. The modified data is calculated by summing particle counts from the smaller high-resolution bins into the larger low-resolution bins. Linear interpolation is used when crossing bin boundaries (if necessary). The noise level in the data is treated as a complete unknown and the singular value cutoff for noise suppression is given a value of unity.

Two data set inversions are described here in some detail. Additional examples are given in Appendix ?? which include the use of a variety of resolutions and number-of-views.

The SVD-based reconstruction of BCO4-610 is shown in Fig. 11. The spatial resolution is 220 by 220 pixels. The grayscale is an uncalibrated linear representation of the object density. The gray level outside of the reconstruction circle (in the four corner regions of the bounding square) represents zero density. Darker grays indicate positive density and lighter grays negative density. Reconstruction artifacts seem to be dominated by (1) azimuthal spatial oscillations of approximate 60-fold symmetry, and (2) shadowing associated with large density discontinuities aligned with view angles. Nevertheless, many of the expected object features are apparent. The major exception is the wedge feature presumably attached to the right side of the cylinder. It is placed in a region of large reconstruction artifacts and its size and orientation are comparable to these artifacts. In this reconstruction it is also difficult to tell whether or not the small interior cones are of zero density or of some small but nonzero density. The same is true for the density of the hemispherical dimple at the left of the cylinder.

The pSVD-based reconstruction of BCO4-610 is shown in Fig. 12. The nonnegativity prior was used for enhancement. The initial SVD reconstruction is that shown in Fig. 11. There are several changes in going from the SVD solution to the pSVD solution. Consider the following list.

1. Regions of negative density are virtually eliminated while remaining consistent with the experimental data. Some negative density artifacts remain near the boundary of the reconstruction circle. This is a consequence of noise in the data.
2. Regions of zero density are well defined. The entire region surrounding the object is now

noted to be of essentially zero density. In addition, the hemispherical dimple and the two interior cones are now noted to be of zero density. In all of these areas the nonnoise-related artifacts are absent. There also appears to be a zero density region near the tip of the dense cone. This appears to be the result of an initial negative density region of the SVD solution related to shadowing. This is likely a falsely identified zero density region.

3. All of the object boundaries (density discontinuities) are sharpened and better defined. This general feature of pSVD solutions allows us, in this case, to identify the hemispherical dimple at left and the wedge at right. It also reveals a truer geometry for the various object boundaries.

4. Shadowing artifacts are nearly completely eliminated, but the azimuthal artifacts are not greatly changed within the object. The presence of these artifacts within the pSVD solution indicates that the null space of the projection operator is very large and the choice of prior is still insufficient to overcome the indeterminacy.

For comparison, the filtered back-projection (FBP) reconstruction using the same data is shown in Fig. 13. It shares most of the problems of the SVD solution without the benefits of the pSVD enhancement. Interestingly, FBP does a better initial reconstruction on the central areas of the figure. This is a natural consequence of the method.

We also performed reconstructions on a the finer grid of $n = m = 500$. For this case, our computer could handle only twelve views. The results of these computations are shown in Figs. 14 through 16. The comparison discussion for the 30 view reconstructions is generally applicable here for the 12 view case. There are some notable exceptions and important comments.

1. The angular reconstruction artifacts have a lower spatial frequency due to the fewer number of views (12 versus 30). The shadowing effect from the dense interior object is still present.

2. Exterior zero density regions are well determined, but interior regions are less well determined relative to the 30 view reconstructions. This is a consequence of greater size of the null space within which we search for correctors.

3. The wedge feature is still captured by the pSVD process although its structure is somewhat different in Figs. 12 and 15.

4. The hemispherical inclusion at the left of the main cylinder is not readily apparent with 12 views. A zero density region is indicated, but the geometry is not obvious.

5. The zero density region at the tip of the dense cone is not present in the 12 view case. This is suggestive that it is an artifact in the 30 view reconstruction.




Figure 11: SVD-based reconstruction of BCO4 section 610. The reconstruction is on the circle interior to a 220 by 220 voxel grid. The data are a reduced resolution ($m = 220, v = 30$) approximation of the full set.




Figure 12: pSVD-based reconstruction of BCO4 section 610 using the nonnegativity prior. The reconstruction is on the circle interior to a 220 by 220 voxel grid. The data are a reduced resolution ($m = 220, v = 30$) approximation of the full set.




Figure 13: Filtered-back-projection reconstruction from the identical data set used in the 30-view reconstructions shown in Figs. 11 and 12.




Figure 14: SVD-based reconstruction of BCO4 section 610. The reconstruction is on the circle interior to a 500 by 500 voxel grid. The data are a reduced resolution ($m = 500, v = 12$) approximation of the full set.




Figure 15: pSVD-based reconstruction of BCO4 section 610 using the nonnegativity prior. The reconstruction is on the circle interior to a 500 by 500 voxel grid. The data are a reduced resolution ($m = 500, v = 12$) approximation of the full set.




Figure 16: Filtered-back-projection reconstruction from the identical data set used in the 12-view reconstructions shown in Figs. 14 and 15.

3.3 Comparative Analysis

The preceding discussions are suggestive that the pSVD is superior to both basic SVD and FBP for the inversion of sparse noisy data. This section provides a simple comparative analysis of the density profiles.

Up until this writing, the geometric and material property details of the BCO4 have remained unknown to us. We chose this blind exercise for two reasons. First, we desired that our reconstructions would remain unbiased by any knowledge of the object other than that gleaned from the data. Second, we are working on other methods of reconstruction enhancements and wish to remain ignorant of the object details until these methods have been applied. Because of this stance, we choose not to perform quantitative comparisons at this time, though we encourage others to do so if they desire. In the near future we will, if possible, report our quantitative findings.

For this report we have chosen to present comparisons of three density profile lineouts of the BCO4 slice 610 for the 30-view reconstruction case. They are shown in Fig. 17. The three lineouts are the voxel columns numbered 100, 120 and 180. For each lineout, we compare the density profiles for reconstructions using FBP (red), SVD (green), and pSVD (blue). The results are shown in Figs. 18 through 20. The density units are chosen so that for each reconstruction the object has a total mass of unity if the reconstruction is taken to be over a unit area.

In general, the pSVD lineouts indicate expected zero-density detection regions, boundary sharpening, and apparent constant density regions. The pSVD solution also shows greater amplitude oscillations relative to the other methods. The FBP lineouts exhibit positive mass out to the fringes of the reconstruction, very poor detection of zero density regions, and a general smoothing of features with no constant density regions indicated.




Figure 17: Lineout locations for density profile comparisons.




Figure 18: Lineout density profiles for column 100 of the BCO4 30-view reconstructions. The three profiles are for FBP (red), SVD (green), and pSVD (blue) reconstructions.




Figure 19: Lineout density profiles for column 120 of the BCO4 30-view reconstructions. The three profiles are for FBP (red), SVD (green), and pSVD (blue) reconstructions.




Figure 20: Lineout density profiles for column 180 of the BCO4 30-view reconstructions. The three profiles are for FBP (red), SVD (green), and pSVD (blue) reconstructions.

4 AHF and the Number-of-Views Question Revisited

One of the key motivations for studying SVD methods for tomographic reconstructions is the design and feasibility of the Advanced Hydrotest Facility (AHF) planned to be constructed at LANL. It is important to quantify the question: Will AHF provide images and reconstructions with sufficient quality of information to justify its existence? SVD and pSVD methods can be used to study certain aspects of this question. This section re-examines the question of how many radiographic views are necessary for a reconstruction of a given quality. This question has been previously visited [ref], but the analysis was done without the benefit of pSVD enhancements.

Full quantitative answers are not yet possible. They require knowledge of the particular metrics of interest. However, the type of study presented here is fully applicable to any metric. And, in many cases, may even give qualitative results one might expect for typical AHF experiments.

Two studies examine the reconstruction quality dependence on the sparsity of data.

4.1 Fixed Resolution Study

This first study uses a variable number of views ($3 \leq L \leq 20$) and fixed resolutions ($K = 100$ and $J = 200$) to examine the relative merit of SVD and pSVD reconstruction methods. Two objects were chosen for this study: TO3 as a discretized density example and TO4 as a smoothly varying density example. The L_2 norm serves as the function of merit for comparison. All reconstructions are performed on noisy data.

Reconstruction samples for TO3 are shown in figure 21. The left column figures show the standard (noise non-amplifying) SVD reconstructions for (a) six, (c) 12, and (e) 20 views. The right column figures show the corresponding pSVD reconstructions obtained using the discretized density prior. Thus figure 21d is the same as figure 6f. A quick study of these reconstructions might convince the reader that the quality of (b) and (c) are similar and also that the quality of (d) is somewhat better than that of (e). This quick assessment is confirmed in figure 22 which shows the norms as a function of number of views and of reconstruction method. Open circles correspond to SVD reconstructions and closed circles to pSVD reconstructions. The letters correspond to the reconstructions in figure 21. Consider the following points.

1. All reconstructions improve with the inclusion of more radiographic views.
2. The pSVD method is always much preferred over the SVD method for a given number of views. This is true if as long as a reasonable prior is chosen.
3. SVD to pSVD improvement can be significant even when the number of views is few.
4. Keep in mind that these norms are on reconstructions obtained from 1% to 5% noisy data. It is reasonable to expect higher or lower norm values depending upon the character of the noise in the data of interest.
5. Reconstruction (f) is exact in spite of noisy data due to the good choice of prior. Note that the SVD reconstruction will not become exact before 50 views (with noiseless data!).
6. In this example the norms of (b) and (c) are similar as noted previously by eye. Also, the

norm of (d) is less than that of (e) as also expected. pSVD reconstructions require about half the number of views as SVD reconstructions to obtain similar figures of merit.

7. In this example for eight or more views, the norms are reduced by at least a factor of two by employing the pSVD enhancement.

Reconstruction samples for TO4 are shown in figure 23. The left column figures show the standard (noise non-amplifying) SVD reconstructions for (a) six, (c) 12, and (e) 20 views. The right column figures show the corresponding pSVD reconstructions obtained using the bounded-density prior. Thus figure 23d is the same as figure 7f. Figure 24 shows the norms as a function of number of views and of reconstruction method. Open circles correspond to SVD reconstructions and closed circles to pSVD reconstructions. The letters correspond to the reconstructions in figure 23. The general comments discussed with TO3 are also applicable here. But a few additional comments are appropriate. The SVD reconstructions for TO4 are of similar merit to those of TO3. But because of the absence of a true object-capturing prior, the pSVD reconstructions for TO4 are of less merit than those of TO3. For this example, pSVD reconstructions require a little more than half the number of views as SVD reconstructions to obtain similar figures of merit, and the SVD to pSVD improvement is generally no better than a factor of two even for many views.




Figure 21: Reconstruction examples of object TO3 with fixed resolution ($K = 100$ and $J = 200$) and variable number of views. The left column figures are SVD reconstructions for six, 12, and 20 views respectively. The right column figures are the corresponding pSVD enhancements using the discretized densities prior. These reconstructions were made from noisy data.




Figure 22: Reconstruction L_2 norms for object TO3 with fixed resolution ($K = 100$ and $J = 200$) and variable number of views. Open circles represent the SVD method and closed circles the pSVD method. The letters correspond to the noisy-data reconstructions of figure 21.




Figure 23: Reconstruction examples of object TO4 with fixed resolution ($K = 100$ and $J = 200$) and variable number of views. The left column figures are SVD reconstructions for six, 12, and 20 views respectively. The right column figures are the corresponding pSVD enhancements using the discretized densities prior. These reconstructions were made from noisy data.




Figure 24: Reconstruction L_2 norms for object TO4 with fixed resolution ($K = 100$ and $J = 200$) and variable number of views. Open circles represent the SVD method and closed circles the pSVD method. The letters correspond to the noisy-data reconstructions of figure 23.

4.2 Fixed Number-of-Views Study

This second study uses a fixed number of views ($L = 12$) and variable resolution ($K = 16, 32, 64, 128,$ and 256 with $J = 2K$) to examine the relative merit of SVD and pSVD reconstruction methods. The same two objects were chosen for this study (TO3 and TO4). The L_2 norm again serves as the function of merit for comparison. All reconstructions are performed on noisy data.

Reconstruction samples for TO3 are shown in Fig. 25. The left column figures show the standard (noise non-amplifying) SVD reconstructions for reconstruction resolutions of (a) 16, (c) 64, and (e) 256. The right column figures show the corresponding pSVD reconstructions obtained using the discretized density prior. The reconstruction merit values for all computed resolutions are shown in Fig. 26. The merit function value is computed by comparing the reconstruction densities to a finely sampled approximation of the actual object ($K = 1024$). Consider the following points.

1. Increased resolution improves the reconstruction.
2. SVD reconstructions have similar merit values at all of the high resolution values. The pSVD reconstructions, however, show a noticeable continued improvement even at the highest resolutions in this study.
3. At the high resolutions, the pSVD reconstructions have merit values nearly twice as good relative to the SVD reconstructions. At the low resolutions, the merits are similarly poor.
4. Note that even a low resolution pSVD reconstruction does a better job of capturing the original object than a high resolution SVD reconstruction. For example, note that reconstruction d is preferred over reconstruction e.

Reconstruction samples for TO4 are shown in Fig. 27. The reconstruction merit values for all computed resolutions are shown in Fig. 28. The interpretation of these figures is similar to that of the previous discussion of TO3. The exception is the use of the nonnegativity prior in computing the pSVD solutions. The conclusions are essentially unchanged.

Note, however, that the reconstruction merits for TO4 are larger than those for TO3. We expect this result because the prior/object match is best for TO3.

What we have found is that for a variety of cases simple priors can be used to significantly reduce the required number of views or required reconstruction resolution given a reconstruction merit constraint.




Figure 25: Reconstruction examples of object TO3 from noisy data, a fixed number of views, $L = 12$ and variable detector and reconstruction space resolutions. The left column figures are SVD reconstructions for $K = 16, 64, 256$, respectively. The right column figures are the corresponding pSVD enhancements using the discretized densities prior. These reconstructions were made from noisy data.




Figure 26: Reconstruction L_2 norms for object TO3 from noisy data, a fixed number of views and variable reconstruction space resolution K . Open circles represent the SVD method and closed circles the pSVD method. The letters correspond to the noisy-data reconstructions of Fig. 25.




Figure 27: Reconstruction examples of object TO4 from noisy data, a fixed number of views, $L = 12$, and variable detector and reconstruction space resolutions. The left column figures are SVD reconstructions for $K = 16, 64, 256$, respectively. The right column figures are the corresponding pSVD enhancements using the discretized densities prior. These reconstructions were made from noisy data.




Figure 28: Reconstruction L_2 norms for object TO4 from noisy data, a fixed number of views and variable reconstruction space resolution K . Open circles represent the SVD method and closed circles the pSVD method. The letters correspond to the noisy-data reconstructions of Fig. 27.

5 Conclusion

Here recap results and discuss extensions.

6 Appendix: pSVD Reconstruction Examples

We include in this appendix additional examples of the pSVD method that are collected from various test cases. Discussion of these examples is limited to providing test conditions and pointing out some interesting features. All of the figures (Figs. 29 through 43) show the object (a) and five reconstructions using the following priors: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.

Figures 29 through 43 show reconstructions of voxelized versions of various objects. Because the objects are exactly defined on a square grid, it is, in principle possible to obtain an exact reconstruction (unlike any examples shown thus far). Often we cannot attain an exact reconstruction because inversions are underdetermined and the data is noisy. The data noise level is 1% to 5%, and the test cases are for $K = 50, J = 100, L = 12$. For simple objects and for good choices of prior the reconstructions are exact in spite of limited noisy data. Some interesting results follow.

Fig. 29: This is a spatially discretized approximation of TO1. This simple object is easily reconstructed exactly with a suitable choice of prior.

Fig. 30: This is a spatially discretized approximation of TO2. This object has more detail than TO1 but it is also reconstructed exactly without difficulty.

Fig. 31: This is a spatially discretized approximation of TO3. The final reconstruction is quite good with only six pixels different from the actual object. The large value for the L_2 norm (0.054) reflects the fact that the six off values differ from the exact value by a good margin (0.25).

Fig. 32: This is a spatially discretized approximation of TO4. This discretization allows a somewhat better pSVD reconstruction (d) than that of Fig. 7 (see Table 1).

Fig. 33: This is a spatially discretized approximation of TO5. As in the previous example, this discretization allows a somewhat better pSVD reconstruction (d) than that of Fig. 8 (see Table 1). But this object is just not well reconstructed using the priors of this report.

Fig. 34: A simple Gaussian is readily reproduced with good accuracy. The initial SVD reconstruction has no nonnegative values; thus, the priors do not improve accuracy.

Fig. 35: This object is similar to TO4 with a central region altered by the addition of a 1/3-size rotated copy of itself.

Fig. 36: This object is a random positive density pattern modulated by a Gaussian. This is the most difficult test case because the object is characteristic of singular vectors associated with very small singular values and of null vectors. The pSVD solutions cannot provide significant improvement mainly due to the lack of negativity corrections as was noted with Fig. 34).

Fig. 37: This object is a small binary-density object within the field of view. It shows the ease with which simple objects are exactly captured with pSVD with noisy data.

Fig. 38: This object is of trinary density with $\rho = 1$ for the inner circle, $\rho = 0.6$ for the ring, and $\rho = 0$ exterior to these. Note that the five-density discretized prior does poorly because the ring density is not in the list of assumed densities.

Fig. 39: This object is a collection of Gaussians superimposed by addition and renormalized to have a maximum density of unity. The reconstructions using simple priors are fairly good. Some of the discrepancy in (c) may be due to the fact that the object is actually defined over the entire square grid while the reconstructions are performed only within the circle interior to this square.

Fig. 40: This object is similar to TO4. It was constructed by thresholding TO4 between densities of 0.1 and 0.9 and then rescaling linearly back to the density interval $[0,1]$. Thus, it has significant regions of both zero and unit density while being smooth elsewhere. The reconstructions are much better than those of TO4.

Fig. 41: This object is a unit density ring with zero interior density and a random selection of zero density voids. This is one of the most dramatic examples of pSVD. Remember that these reconstructions are from noisy data.

Fig. 42: This object is a binary-density version of TO4 using a density threshold of 0.3.

Fig. 43: This object is a unit density circle with a simulated-fracture interior void pattern. The pSVD reconstructions are stunningly accurate.




Figure 29: Reconstruction examples of a voxelized version of object TO1. This object (a) has binary density and is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 30: Reconstruction examples of a voxelized version of object TO2. This object (a) has binary density and is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 31: Reconstruction examples of a voxelized version of object TO3. This object (a) is of discrete density and is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 32: Reconstruction examples of a voxelized version of object TO4. This object (a) is a discrete density approximation to a smooth variation and is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 33: Reconstruction examples of a voxelized version of object TO5. This object (a) is a discrete density approximation to a smooth variation and is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 34: Reconstruction examples of a discretized simple Gaussian. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 35: Reconstruction examples of an object similar to TO4. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 36: Reconstruction examples of a random nonnegative density pattern modulated by a Gaussian. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 37: Reconstruction examples of a small unit density object. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 38: Reconstruction examples of a trinary-density object with $\rho = 1$ for the inner circle, $\rho = 0.6$ for the ring, and $\rho = 0$ exterior to these. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 39: Reconstruction examples of a sum of randomly-placed and sized Gaussians. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 40: Reconstruction examples of a modified TO4 (see text for the exact object description). This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 41: Reconstruction examples of a unit-density ring with random zero-density voids. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 42: Reconstruction examples of a binarized TO4 (density threshold of 0.3). This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.




Figure 43: Reconstruction examples of a unit density circle with a simulated-fracture interior void pattern. This object (a) is exactly described on a 50 by 50 grid of densities. The pSVD reconstructions use the following prior knowledge: (b) none; (c) $\rho \geq 0$; (d) $0 \leq \rho \leq 1$; (e) $\rho \in \{0, 1\}$; (f) $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The L_2 merit function value is printed next to each reconstruction.

7 References